

ISIS - Question #1141

cam2map very slow in cluster environment since Isis 3.4.0

2012-09-27 01:20 PM - Tammy Becker

Status:	Closed
Priority:	Normal
Assignee:	Jeff Anderson
Category:	Applications
Software Version:	
Description	
External Post: https://isis.astrokeology.usgs.gov/IsisSupport/index.php/topic.3567.0.html	
Dear ISIS team,	
we are running a CentOS 5 cluster with a self-compiled ISIS version (since you are not providing the libraries for RHEL5 any more). We realize problems running cam2map on the cluster if there is heavy load on the file system.	
Running the older version 3.3.1 makes no problems, also running the same binary (3.4.1, also 3.4.0 tested) on a local scratch disk makes no problems at all.	
The file system is xfs mounted via nfs, which is also mounted to a bunch of other cluster members with some load on the file system (around 100 processes). When I run strace with the cam2map command, these are my last lines:	
Code:	
<pre>open("G02_018928_1717_XN_08S082W.lev1eo.cub", O_RDONLY) = 10 lseek(10, 391075732, SEEK_SET) = 391075732 read(10, "Object = mroctx2isis\n IsisVersi"..., 8191) = 2390 close(10) = 0 open("/isis-3.4.1/version", O_RDONLY) = 10 read(10, "3.4.1.4648\n2012-08-17 # Versio"..., 8191) = 135 close(10) = 0 open("G02_018928_1717_XN_08S082W.lev2.cub", O_RDWR) = 10 lseek(10, 0, SEEK_END) = 731533740 lseek(10, 0, SEEK_CUR) = 731533740 lseek(10, 0, SEEK_CUR) = 731533740 lseek(10, 0, SEEK_END) = 731533740 lseek(10, 0, SEEK_CUR) = 731533740 lseek(10, 731533740, SEEK_SET) = 731533740 writev(10, [{NULL, 0}, {"Object = mroctx2isis\n IsisVersi"..., 3114}], 2) = 3114 close(10) = 0 fstat(1, {st_mode=S_IFCHR 0620, st_rdev=makedev(136, 0), ...}) = 0 mmap(NULL, 4096, PROT_READ PROT_WRITE, MAP_PRIVATE MAP_ANONYMOUS, -1, 0) = 0x2adca70ca000 write(1, "cam2map: Working\n", 17cam2map: Working) = 17) = 13ocessed\r", 130% Processed lseek(8, 1180237824, SEEK_SET) = 1180237824 read(8, "z\20y\20x\20w\20s\20r\20l\20s\20q\20n\20l\20k\20j\20i\20g\20"..., 32768) = 32768 futexp(0x3ab300d350, FUTEX_WAKE_PRIVATE, 2147483647) = 0 lseek(8, 1180270592, SEEK_SET) = 1180270592 read(8, "\275\17\273\17\271\17\267\17\265\17\263\17\262\17\263\17\252\17\255\17\237\17\341\16\321\r\332\f\302\f\3\r"..., 32768) = 32768</pre>	
Do you have any clue what could be the reason for the break down?	
Thank you very much! Sebastian	

History

#1 - 2012-10-01 12:33 PM - Tammy Becker

Jeff, I've assigned this post to you since you have IT looking into it and you've responded to the external post.

Not sure who's paying for your time on this??

#2 - 2012-10-02 08:43 AM - Jeff Anderson

Response written on support board

We run our own cluster and NFS can be a serious problem depending on the number nodes and whether the programs are I/O intensive or CPU intensive.

You might check your NFS server and how it is responding to requests from the cluster. Is it getting a higher number of NFS read/write requests with the new version of cam2map versus the old version? Here is why I could see cam2map impacting your cluster now versus in the past. There were recent changes such that the new version is much faster nearly two to three times. That means cam2map doesn't need as much CPU time and needs to get data off disk faster. When put on a cluster all the instances of cam2map are hammering the NFS server. Now your cluster isn't number crunching as much anymore ... it's waiting for I/O requests to be completed.

We have two tiers of storage. For I/O intensive applications we are using a Lustre cluster with 40Gbit/s infiniband networking. It is very fast and we can get throughput of nearly 2GBytes/second and can grow that by adding extra lustre nodes.

We also use NFS also backed by 40Gbit/s infiniband. However, we spread that I/O over six NFS servers. If we only had one NFS server it would get bogged down. The work load is spread out by power users. That is, our heaviest ISIS users don't have data on the same NFS servers.

Hope this helps.

#3 - 2012-10-15 09:13 AM - Jeff Anderson

Closed ticket in support board.

#4 - 2012-10-29 10:37 AM - Tammy Becker

Reporter never responded, consider the issue closed;